

A Generic Programming Model
for
Network Processors

2nd Year Report

Kevin Lee

*Supervisors: Dr. Geoff Coulson and
Prof. Gordon Blair*

Computing Department
Lancaster University

September 2004

Contents

1	Introduction	3
2	Motivation	3
3	Our Approach	4
4	Progress Report	5
4.1	First Year	5
4.2	Second Year	5
4.3	Publications	5
5	Plan for Final Year	7
5.1	Programming Model Implementation	7
5.2	Component Framework Implementation	8
5.3	Analysis and Evaluation	8
5.4	Phase IV: Writing Up	9
5.5	Ongoing tasks	10
6	Concluding Remarks	11
7	Appendix A: International Conference on Networks (ICON) 2004 Paper	13

1 Introduction

This document provides an overview of the research that has been carried out as part of the author's Ph.D. in the Department of Computing, Lancaster University. In particular, it focuses on the progress achieved over the last twelve months since the last assessment (*“End of First Year Report”*). The report outlines the motivation for the research, as well as the principles and the main ideas of the approach. It also briefly outlines the progress achieved so far and discusses the plan for the final year of the project.

2 Motivation

Network Processors (NPs) are an attempt by hardware vendors to fulfill the growing need for low-priced specialized network hardware elements that are more future proof than conventional custom hardware or Application Specific Integrated Circuit (ASIC) based designs, and can be applied in a wider range of situations (e.g. in networked devices, as edge-network routers and even in the network core). In addition, NPs are seen by some as potential vehicles for the deployment of active networking-derived technologies [1] which exploit the potential of NPs for run-time software reconfiguration. Architecturally, NPs are multiprocessor-based hardware units that support a number of network ports and provide software-programmable packet processing facilities. They have the ability to perform relatively complex packet processing at line speeds.

There is a downside to current NP designs, however: they are notoriously difficult to program [2], [3]. This is because of their complex design (e.g. involving multiple processors, including both general purpose and specialized processors; and multiple memory and interconnect technologies), their extreme architectural heterogeneity across vendors and products [4], and their demanding performance constraints.

In addition to high con-figurability, it is recognized [5] that a desirable characteristic of NPs is runtime reconfiguration. This characteristic is useful in a NP environment for a number of applications, including adaptive security provision [6], dynamically extensible services (PromethOS) [7], and even in space-based applications [8]. In addition the active networking (AN) community have been heavily involved in investigating the use of NPs in their field [9]. This is because ANs require significantly more data-plane processing and also require routers to expose their state to allow reconfiguration of forwarding functions.

Furthermore, NPs which exhibit richly-featured hardware designs largely remain under-exploited by software [10]; and their extreme heterogeneity tends to inhibit translation of software, software designs, or even skills across brands. The problem is exacerbated by the need for extreme performance and high con-figurability, both of which add considerably to software complexity.

Unfortunately because of the complexity of NPs, current NP software toolkits fail to provide any support at all for runtime reconfiguration. The majority of NP software toolkits have focused on enhancing the ability of compilers [11]

and intelligent code placement [12]. Runtime support for configuring and reconfiguring NPs is extremely basic and for the majority of cases consists of halting a platform, uploading new code blocks and restarting the platform.

To solve these problems a generic NP software toolkit would be required to support design portability across heterogeneous NP architectures and additionally support runtime reconfiguration. This must, however take account of the overriding aim being to do this whilst maintaining the extreme-performance that NPs are capable.

3 Our Approach

The aim of the research is to solve the problem of NP heterogeneity and aid in easing the difficulty of developing software solutions for NP platforms. In addition to supporting the array of heterogeneous NP architectures currently in existence, the research is also attempting to support future NP architectures.

Our goals are to develop a generic programming model for NPs that accommodates complex architectures and architectural heterogeneity while also supporting design portability, high performance and runtime reconfigurability. Our approach is based on the run-time component model presented in [13]. It features a component runtime with low memory footprint, employs formally specified interfaces, supports components written in different programming languages, and uniformly abstracts over different processor types and different inter-processor communication mechanisms without loss of performance. In addition, the programming model utilises the notion of component frameworks [14] to constrain sets of components within an area of encapsulation.

The model promotes conceptual uniformity and design portability across a wide variety of NP types while simultaneously exploiting hardware assists that are specific to individual NPs. The programming model's architecture exhibits built-in support for third-party loading and binding of components. This is accomplished by the use of pluggable loaders and binders which load binary code and can bind binary code components together. This is especially relevant for the deployment of components on NP platforms as loading and binding inherently occurs in a third-party manner. Furthermore, because of the nature of the platform, and the fine granularity necessary to perform reconfiguration successfully, it is necessary to support binding in many different ways, the programming model supports this.

As well as dealing with the wide variety of heterogeneous NP architectures, the model, because of its inherent modular design and high-level support also lends itself to supporting runtime reconfiguration. This high-level support takes the form of resource management and CF composition both vital for supporting effectively a runtime reconfigurable platform.

4 Progress Report

In this chapter I present a brief summary of the work I undertook in my first and second years. I also present a list of my publications with brief summaries.

4.1 First Year

My first year involved a broad investigation of the Active Networking field, this investigation was spawned from the target field of the NetKit [15] research project. I was also heavily involved in the redesign and porting of the MSCOM based OpenCOM to the XPCOM version of OpenCOM on the linux platform.

The bulk of my first year was involved in the general task of enhancing my programming and operating system skills which have since proved to be vital. This involved extensive C/C++ programming, including writing a IDL compiler for the redesigned OpenCOM. I also greatly improved my Linux kernel/system/userspace programming knowledge which has proved vital in setting up and fully investigating the IXP platforms.

4.2 Second Year

Further to the work porting OpenCOM to the linux environment, i was Involved in the further redesign and reimplementaion of OpenCOM to a smaller footprint minus the dependency of XPCOM.

We have been working to deploy and evaluate the OpenCOM component model on the Intel IXP1200. The IXP1200 was selected because of its open and well documented architecture, and because it is a richly-featured NP in terms of the dimensions discussed in [13]. In addition to the study and implementation work on the IXP1200 platform, I carried out a extensive initial survey of the field of NPs.

4.3 Publications

The following list describes published papers that present partial results of our research. Kevin Lee has been the primary author of the papers cited below.

- “*Towards a Generic Programming Model for Network Processors*” (Kevin Lee, Geoff Coulson, Gordon Blair, Ackbar Joolia, Jó Ueyama) IEEE International Conference On Networks ICON 2004. This paper firstly addresses the difficulty to program several existing networking processor-based routers (like IXP1200 and Motorola [16]). As a solution, we propose a generic programming model to facilitate programming of these complex environments.

In addition, the following papers have been published by our research group (NetKit project). I has been actively involved in writing these papers, and indeed has contributed large sections to each paper:

- “*NETKIT: A Software Component-Based Approach to Programmable Networking*” (Geoff Coulson, Gordon Blair, David Hutchison, Ackbar Joolia, Kevin Lee, Jó Ueyama, Antônio Gomes, Yimin Ye) appeared in *Computer Communication Review*, a publication of ACM SIGCOMM, Volume 33, Number 5, October 2003. ISSN # 0146-4833. This journal paper describes the approach of our research and presents partial results collected from our implementation.
- “*A Reflective Middleware-based Approach to Programmable Networking*” (Geoff Coulson, Gordon Blair, Antônio T. Gomes, Ackbar Joolia, Kevin Lee, Jó Ueyama, Yimin Ye), at the Proceedings of the 2nd International Workshop on Reflective and Adaptive Middleware (located with ACM/IFIP/USENIX Middleware 2003), Rio de Janeiro, Brazil, June, 2003. This paper presents how our component model (OpenCOM), which was successfully applied on middleware platforms, can be implemented to enhance the flexibility on programmable networks.
- “*A Re-configurable Component-based Model for Programmable Networks*” (Jó Ueyama, Stefan Schmid, Geoff Coulson, Gordon Blair, Antônio Tadeu Gomes, Ackbar Joolia, and Kevin Lee), presented at the Fourth International Workshop on Distributed Auto-adaptive and Reconfigurable Systems (DARES 2004), in conjunction with ICDCS2004, Tokyo, Japan, March 23-26, 2004. The current paper describes how our component model can better support the configuration and reconfiguration mechanisms. We then show the study case applied to networking systems.
- “*A Globally-Applied Component Framework for Programmable Routers*” (Jó Ueyama, Geoff Coulson, Gordon Blair, Stefan Schmid, Antônio T. Gomes, Ackbar Joolia, Kevin Lee), presented at the Fifth International Working Conference on Active Networks IWAN 2003, 10 - 12 December 2003, Kyoto, Japan. This paper describes the main principles of our research (namely our ‘globally-applied’ approach for programmable networking systems). We demonstrate the advantageous of this approach (e.g. we describe that the programming of these complex environments, like the IXP1200, can be simplified by providing a unique component model that can deploy Linux and microengine components). At last it gives some details of the implementation being carried out.
- “*Enabling Re-configurability on Component-based Programmable Nodes*” (Jó Ueyama, Stefan Schmid, Geoff Coulson, Gordon S. Blair, Antônio T. Gomes, Ackbar Joolia, Kevin Lee), PhD Workshop, Combined DAIS/FMOODS conference, Paris, Nov 2003. This short paper served as a survey of our research and presents the first results of our work.

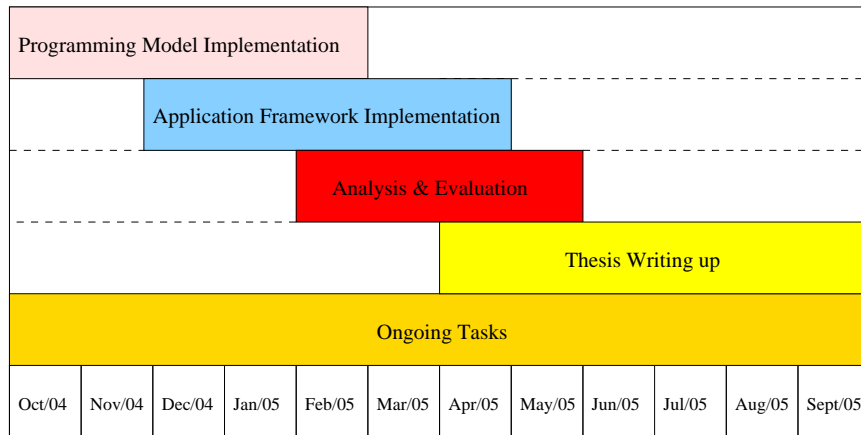


Figure 1: Timetable for 2004/2005

5 Plan for Final Year

This years workload will be broadly split up into five tasks; implementation of the programming model, implementation of a component framework, analysis and evaluation, thesis writeup and ongoing Tasks (see below). This chapter details each of these and presents a estimated timeline for each of the tasks. The diagram in figure 1 provides a estimated timeline for the project over the next 12 months.

5.1 Programming Model Implementation

This phase comprises of the implementation of the programming model for the Intel IXP2400, the core IP implementation, example application and testing.

I am currently about half way though the implementation of our programming model for the Intel IXP platforms. We have a mature component-model which can run on any gcc-supported general purpose processor. We also have this deployed successfully on the Strong-ARM processor of the IXP1200. The remainder of the implementation involves polishing the deployment of the component model on the IXP, implementation of more mature binders to complement the loaders already implemented.

I am currently in the process of transferring and upgrading my skills from the IXP1200 [17] platform to the IXP2400 platform [18]. The reason for this is the improved tool-set support as well as the availability of modular IPv4 and (possibly) TCP termination source-code. The availability of this source code will greatly speed-up the implementation process.

A certain degree of genericity can be demonstrated by implementing the programming model on both the IXP1200 and IXP2400 due to the substantial differences in the architectures, the IXP1200 was originally designed by NEC and re-engineered by Intel into the IXP2400.

However I am aware that this is not a powerful enough study alone to illustrate true genericity. Thus it is my intention to investigate the possibility of deploying the platform on another NP. As such i am actively researching other NPs with the aim of acquiring a substantially different architecture to deploy the programming model. Potential NPs include the AMCC nP3710 [19] and the Agere PayLoadPlus [3].

It is expected that this phase will be completed within 5 months, although this may depend on any modifications needed for the analysis phase. In addition, any additional platform implementation will take additional time.

5.2 Component Framework Implementation

To demonstrate the power of the approach introduced in section 3, as well as the programming model implementation itself a component framework will be developed. This will consist of a domain-specific framework which allows the plugging in of application-specific components. This component framework will be used to analyse and evaluate the performance of the programming model.

NPs are designed for data processing a large number of packet flows at line speeds, as such a video-on-demand (VoD) application is the perfect evaluation application. An component framework for supporting VoD would enable a NP to deal with a large number of video streams targetted toward specific client demands. NPs have the ability to not only route a large amount of streams to a wide variety of clients, but also the processing power to adapt the streams (the IXP2400 for instance has around 10Ghz of processing power, for 4 gigabit ports). This adaptation can take the form of transcoding, B-frame dropping or wavelet selecting.

The actual implementation will consists of a component framework based on the programming model implementation. A number of implemented video codecs for the IXPs microengines will also be developed to test the frameworks performance. These will be heavily used during the analysis and evaluation phases.

5.3 Analysis and Evaluation

After the implementation of the programming model and application framework, it will be necessary to perform a extensive and detailed analysis and evaluation stage. The focus of these tasks is to investigate the implementation in relation to the thesis discussed in [13] and will be targeted towards supporting the writing of further publications as well as the thesis document.

The analysis and evaluation will consist broadly of two groups, a qualitative group and a quantatative group. The qualitative group of tests will consist of tests aimed at proving the value-added aspects of the programming model are sound. For instance a number of test-cases will have to be carried out to study the act of runtime reconfiguration of the application framework. The qualiatitive tests will also investigate how the mechanisms used differ from the

ones utilised by other programming models (like the Intel ACE [20] framework) and evaluation any benefits gained.

The quantitative group of tests have the broad aim of analysing the performance of the implemented programming model and application framework. It will utilise the implemented VoD scaling/degrading plugins in a application scenario to show how system performance is affected when processors handling gigabit speed traffic is reconfigured.

As well as evaluating the implementation in a qualitative and quantitative manner, as the programming model is claiming to be generic, it is necessary to evaluation on multiple heterogeneous architectures. As mentioned in section 5.1 the aim is to evaluate the programming model with the component framework on the IXP1200, IXP2400 and another NP. Thus, the evaluation of the model on the can utilise the StrongARM (IXP1200), XScale ARM (IXP2400), MicroEnginesV1 (IXP1200), MicroEnginesV2 (IXP2400) as well as bindings for coprocessors on these platforms. In addition to these, the component framework will be implemented on a x86 linux-based architecture with our current programming model implementation.

5.4 Phase IV: Writing Up

The final deliverable of this research will be a thesis document. This is expected to take approximately nine months from early drafts to the final document. A preliminary plan for the structure of this document together with comments is provided below:

1. **Introduction** This section will introduce the reader to the research topic and the motivation for my work. It will introduce the field on NPs, and try to sum up the important peculiarities of that field. It will first discuss the positive attributes NP can bring to the field of active networks and then introduce the problems with current NP software. The remainder of the introduction will involve the motivation for the work, giving a detailed introduction to the proposed solution. This section will also contains a brief description of the structure of the thesis itself.
2. **State of the Art of Network Processors** This will section outline the current state of the art in the field of NPs. Details important distinctive architectures, provides a classification of these architectures and discusses important development software.
3. **State of the Art of Component-based Network Processor Software** This section will outline important milestone research efforts in the field of component-based systems software.
4. **Generic Programming Model** This section details the reasoning behind the notion of a generic programming model together with a detailed look at the model itself. It will specifically focus on issues of utilizing a component-based programming model in the field of networking and NPs.

5. **Intel IXP2400 Implementation** This section will detail the implementation of the generic programming model for a single target NP architecture. The detail will include the core model as well as a IP 13 forwarder implementation and a example active application like a video transcoder.
6. **Generic Application to NPs** This section will build on the previous two chapters by illustrating in detail how the programming model maps in theory and in practice to a range of NP architectures. The aims of this chapter are to show that all the distinct architectures previously identified in the NP classification are supported in the generic programming model.
7. **Evaluation** This section will include an evaluation of the implementation and analysis undertaken utilizing the IXP2400. Importantly, this section will also evaluate the notion of a generic programming model and the issues this presents.
8. **Conclusions and Future Work** This section will give a detailed sum-up of the research presented in the thesis and present conclusions based on this. Part of this section will also be devoted to discussing potential future work based on this research.

5.5 Ongoing tasks

In addition to the focused tasks of implementation and evaluation, an equally important part of the subsequent year will two ongoing tasks. These are:

- Ongoing Research in the field on NPs

As the premise of the thesis is that the described programming model for NPs has the property being generic across multiple platforms, it is essential to illustrate this in the thesis. This will take the form of a investigation into the entire field of NPs. In order to prove the generic aspects of the thesis, a number of distinct NPs identified as being samples of the different identified NP categories will be analyzed in detail. If time is available, the ideal is to investigate the sample NPs programmably including a investigative implementation. We are already actively attempting to obtain NP products other than those donated to us from Intel. The result of this research will directly be applicable on potential publications and the thesis document itself.

- Writing/Publishing further papers

The writing and submitting of further papers is an important ongoing task in the next year. The first phase of publications included the ICON 2004 paper present in Appendix A. This paper represented medium to mature level project position and also early implementation results. The second phase of publications will depend on the completion of the implementation phase and will involve presenting this data heavily. It is hoped that this phase will consist of a single conference/workshop paper which will be

written in the weeks after implementation/evaluation is completed. A third phase is also expected towards the end of the thesis writeup when the complete refined thesis has been formulated and fully researched. This phase would entail a final paper presenting the completed implementation together from the final thesis and backup research. It is expected that this would be targeted towards a journal because of the expected length of such a paper.

6 Concluding Remarks

This report has presented a brief introduction to the field, an assessment of the perceived problem and the motivation for the research. It has also presented a summary of the research conducted in the first and second year, both personally and as part of the research group. A plan of work for the next 12 months has been presented together with a rationale for this work. It is hoped that this plan will represent the work for the next 12 months, however it is inevitable that minor changes will be necessary.

References

- [1] L. Ruf, R. Pletka, P. Erni, and P. Droz. Towards High-Performance Active Networking. *Proceedings of the Fifth Annual International Working Conference on Active Networks (IWAN 2003)*, December 2003. Kyoto, Japan.
- [2] C. Kulkarni, M. Gries, C. Sauer, and K. Keutzer. Programming Challenges in Network Processor Deployment. In *Int. Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES)*, October 2003.
- [3] Agere Systems Proprietary. The Challenge for Next Generation Network Processors. April 2001.
- [4] M. Tsai, C. Kulkarni, C. Sauer, N. Shah, and K. Keutzer. A benchmarking methodology for network processors. In *1st Workshop on network processors along with HPCA 2002*, February 2002.
- [5] I.A. Troxel, A.D. George, and S. Oral. Design and analysis of a dynamically reconfigurable network processor.
- [6] Scott J. Harper. A secure adaptive network processor.
- [7] L. Ruf, R. Keller, and B. Plattner. A Scalable High-performance Router Platform Supporting Dynamic Service Extensibility On Network and Host Processors. In *Proc. of 2004 ACS/IEEE Int. Conf. on Pervasive Services (ICPS), Beirut, Lebanon*. IEEE, Jul. 2004.
- [8] A. Gray, J. Srinivasan, Y. Chong, A. Farrington, K. Peters, and V. Stanton. Software/hardware reconfigurable network processor for space networks.
- [9] Andreas Kind, Roman Pletka, and Marcel Waldvogel. The role of network processors in active networks. In *Proceedings of IWAN 2003*, pages 18–29, Kyoto, Japan, December 2003.
- [10] P. G. Paulin. Network Processors: A Perspective on Market Requirements, Processor Architectures and Embedded S/W Tools. *STMicroelectronics*, 2001.
- [11] J. Wagner and R. Leupers. C Compiler Design for an Industrial Network Processor. *Proceedings of the 2001 ACM SIGPLAN Workshop on Optimization of Middleware and Distributed Systems*, 2001.
- [12] Intel Corporation. Introduction to the auto-partitioning programming model.
- [13] Kevin Lee, Geoff Coulson, Gordon Blair, Ackbar Joolia, and Jo Ueyama. Towards a generic programming model for network processors. In *to appear in Proc. IEEE International Conference on Networks (ICON.04)*, Singapore, November 2004.

- [14] C. Szyperski. *Component Software: Beyond Object Oriented Programming*. Addison-Wesley, first edition, 1998.
- [15] Geoff Coulson, Gordon Blair, David Hutchison, Ackbar Joolia, Kevin Lee, Jo Ueyama, Antonio Gomes, and Yimin Ye. NETKIT: A Software Component-Based Approach to Programmable Networking . In *ACM SIGCOMM Computer Communication Review*, volume 33, No 5, October 2003.
- [16] Motorola Research. Architecture guide, C-5e/C-3e Network Processor, Silicon Revision B0, 2003.
- [17] Radisys Corporation. IXP1200 White Paper: Using the Intel IXP1200 Network Processor to Optimize Packet-Processing Application Development, 2001.
- [18] M. Adiletta et al. The Next Generation of Intel Network Processors. *Intel Technology Journal, Volume 6, issue 3*, August 15 2002.
- [19] Applied Micro Circuits Corporation. np3700, 5-gbps network processor with integrated traffic manager, product brief, 2003.
- [20] Intel Press. MicroACE, Design Document, revision 1.0. *Intel Press, Intel Corporation*, 2001.

7 Appendix A: International Conference on Networks (ICON) 2004 Paper

TOWARDS A GENERIC PROGRAMMING MODEL FOR NETWORK PROCESSORS

Kevin Lee, Geoff Coulson, Gordon Blair, Ackbar Joolia, Jo Ueyama

Lancaster University, Lancaster, LA1 4YR, UK

Abstract—Network Processors (NPs) are emerging as a cost effective network element technology that can be more readily updated and evolved than custom hardware or ASIC-based designs. Moreover, NPs promise support for run-time reconfiguration of low-level networking software. However, it is notoriously difficult to develop software for NPs because of their complex design, architectural heterogeneity, and demanding performance constraints. In this paper we present a runtime component-based approach to programming NPs. The approach promotes conceptual uniformity and design portability across a wide variety of NP types while simultaneously exploiting hardware assists that are specific to individual NPs. To show how our approach can be applied in a wide range of types of NPs we characterise the design space of NPs and demonstrate the applicability of our concepts to the various classes identified. Then, as a detailed case study, we focus on programming the Intel IXP1200 NP. This demonstrates that our approach can be effectively applied, e.g. in terms of performance, in a demanding real-world NP environment.

I. INTRODUCTION

Network Processors (NPs) are an attempt by hardware vendors to fulfill the growing need for low-priced specialised network hardware elements that are more future proof than conventional custom hardware or ASIC-based designs, and can be applied in a wider range of situations (e.g. in networked devices, as edge-network routers and even in the network core). In addition, NPs are seen by some as potential vehicles for the deployment of active networking-derived technologies [1] which exploit the potential of NPs for run-time software reconfiguration. Architecturally, NPs are multiprocessor-based hardware units that support a number of network ports and provide software-programmable packet processing facilities. They have the ability to perform relatively complex packet processing at line speeds.

There is a downside to current NP designs, however: they are notoriously difficult to program [2], [3]. This is because of their complex design (e.g. involving multiple processors, including both general purpose and specialised processors; and multiple memory and interconnect technologies), their extreme architectural heterogeneity across vendors and products [4], and their demanding performance constraints.

Therefore, NPs often exhibit richly-featured hardware designs that remain underexploited by software [5]; and their extreme heterogeneity tends to inhibit translation of software, software designs, or even skills across brands. The problem is exacerbated by the need for high performance and runtime reconfiguration, both of which add considerably to software

complexity. In particular, because of their complexity, many NP software toolkits fail to provide any support at all for runtime reconfiguration.

The aim of the research discussed in this paper is to develop a generic programming model for NPs that accommodates complex architectures and architectural heterogeneity while also supporting design portability, high performance and runtime reconfigurability. Our approach is based on a runtime software component model. This promotes conceptual uniformity and design portability across a wide variety of NP types while simultaneously exploiting hardware assists that are specific to individual NPs. It features a distributed runtime with low memory footprint, employs formally specified interfaces, supports components written in different programming languages, and uniformly abstracts over different processor types and different inter-processor communication mechanisms without loss of performance. It also explicitly supports run-time reconfiguration of software.

The remainder of the paper is structured as follows. In section II, we characterise the design space of NPs as a basis for arguing the genericity of our approach, and also survey a number of existing programming models provided both by the manufacturers of various NP products, and by independent researchers. In section III, we present our approach to programming NPs and show how this improves on existing approaches. Then, in section IV, we provide a detailed case study of the application of our approach to the Intel IXP1200 NP. Finally, in section V we offer our conclusions.

II. NETWORK PROCESSORS

A. Classification

As mentioned, the field of NPs is notable for its great architectural heterogeneity. In general, however, it can safely be said that NPs universally provide programmable support for processing packets, and that this usually takes the form of one or more *packet processors*. These can be supported either on a single chip or across multiple chips. In addition, NPs universally support a number of MAC-level ports, some memory, and some form or forms of processor interconnect.

In this section we attempt to capture the design space of NPs in terms of a small number of orthogonal dimensions. In particular, we have chosen *four* key dimensions which, we believe, most usefully partition the NP design space. These are:

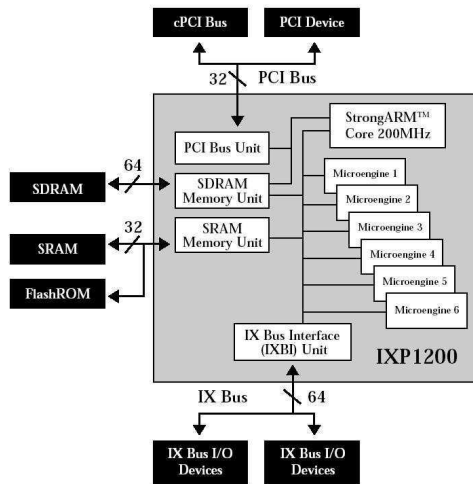


Fig. 1. The Intel IXP1200 (from [7])

- the *packet processor dimension* - the range of types of packet processors supported by an architecture
- the *memory architecture dimension* - the range of technologies and organisations of the memory provided
- the *interconnect dimension* - the range of interconnect technologies employed
- the *control and management dimension* - the degree of support for centralised control and management

We also demonstrate how some prominent NP products map to this space. In so doing, we lay the groundwork for a discussion on how our component-based programming approach can accommodate the full diversity of NPs.

1) *The Packet Processor Dimension:* Most NPs feature multiple packet processors, but the nature of these can vary from CPUs with very general instruction sets to single-purpose dedicated units for, e.g., checksumming or hashing, which are not programmable. Furthermore, some NPs feature only one type of packet processor and others support a number of different types.

For example, the Intel IXP1200 NP [6] (see figure 1) supports a uniform set of six so-called *microengines* which serve as packet processors. These are 233-600Mhz CPUs whose instruction set includes I/O to/from MAC-ports, packet queuing support, and checksumming. They support hardware threads with zero context switch overhead and can be programmed either in assembler or C. The IXP1200 also includes a general purpose StrongARM CPU which serves as a controller and also typically performs slow-path operations.

On the other hand, the Motorola C-Port [8] employs so-called *channel processors* which are generic packet processors grouped in sets of four that share an area of fast memory. But in addition it supports a range of dedicated, non programmable, processors that perform functions such as queue management, table lookup, and buffer management.

As a third example, the EZChip NP-1 [9] has no fully generic processors. Rather, it employs dedicated packet processors that perform specific tasks such as parsing packets,

table lookup or packet modification. Although these are dedicated to their given ‘domain’, they are quite flexible and programmable within that domain.

2) *The Memory Dimension:* Memory is used in all the fundamental operations of a NP, including packet storage, table lookup, queuing and synchronisation. The properties of different memory types typically differ in terms of size and speed, whereas their organisation differs in terms of the degree of centralisation employed and the accessibility from different packet processors.

Memory types and organisations greatly affect the structure of NP software. To deal with the memory organisation of a particular platform, the programmer has to choose the best memory use strategy for a particular operation. For example, when creating a flow-table for high-speed connections an Intel IXP1200 programmer might choose on-chip scratch memory, whereas an IBM PowerNP programmer [10] might use that architecture’s high-speed internal SRAM.

3) *The Interconnect Dimension:* Different NPs provide different mechanisms for inter-processor communication such as shared registers, buses (of varying types), shared memory (perhaps a range of types that make different trade-offs between capacity and speed), and dedicated channels.

For example, the IXP1200 provides a fast bus for communication between its microengines, MAC ports and memory. It also provides shared registers and a range of memory types (i.e. SRAM, SDRAM). The shared registers and memory are typically used together at the software level to realise inter-processor communication. The newer IXP2400 NP from Intel also provides ‘next-neighbour’ registers that provide a dedicated interconnect between two ‘adjacent’ microengines.

The Motorola C-Port employs shared fast memory for interconnection between grouped channel processors (as mentioned above). It also employs multiple onboard buses for communication between these groups, and shared memory that is managed by a dedicated processor.

Unlike the two examples above, the EZChip offers a very static and limited interconnect which arranges the packet processors in a strict pipeline topology. The Cisco PXF [11] uses a variant of this approach: it offers multiple parallel pipelines and some capability for communication between pipelines. Clearly, these architectures are less flexible, although potentially faster, than the bus-based interconnects discussed above.

4) *The Control and Management Dimension:* Apart from the genericity/specificity of their packet processors, different NPs make different choices regarding centralisation/ decentralisation of control and management. For example, some NPs rely exclusively on external control in the form of a host workstation. Others (e.g. the IXP1200) incorporate a commodity CPU on the NP itself which runs an operating system, and others support sufficiently powerful and general packet processors that any of these can potentially serve as a locus of control and management.

The IXP 1200’s on-board StrongARM CPU runs a commodity OS such as Linux. As well as handling slow-path packet processing, the StrongARM is responsible for loading

code onto the microengines and stopping and starting them as required.

The Motorola C-Port, on the other hand, has no built-in centralised controller. Instead, it relies on a host workstation to load and supervise the operation of its ‘channel controller’ packet processors. Nevertheless, it is theoretically possible to dedicate one of the channel controllers to take the supervisory role, especially if fine-grained dynamic reconfiguration of the NP is a goal.

Similarly, the EZChip relies on a host workstation for control and management. In this case, there is no alternative because dedicating one of the packet processors, even if possible (cf. their lack of generality), would introduce an unacceptable bottleneck in the pipeline.

B. Software for Network Processors

The provision of software development environments for different NPs is almost as diverse as NP hardware architecture. In this section we examine both proprietary and research-derived programming environments and show that each is hard to generalise beyond the specific architecture at which it is targeted.

In terms of proprietary software, we focus on programming models and development environments for the IXP1200 and the IBM PowerNP. Information on the software environments used by other NPs is unfortunately hard to obtain without signing non-disclosure agreements.

Intel’s *MicroACE* [12] is targeted at the IXP1200 and other Intel IXA products. In this model, proxy-like software elements (called *active computing elements* or ACEs) on the IXP1200’s StrongARM control processor are ‘mirrored’ by blocks of code (called microblocks) that run on microengines. Thanks to this mirroring, when the programmer loads a StrongARM element, the corresponding microblock is transparently loaded onto a microengine as a side effect. The microblock can choose to offload packets to its associated ACE for handling on the slow path.

Although it provides a useful degree of abstraction, the *MicroACE* approach is limited to IXP1200-like architectures that employ a tightly integrated control processor. Furthermore, the model leaves linkages between microblocks implicit in the way the microblocks are written: is not possible to combine microblocks in unanticipated topologies or to exploit interconnect mechanisms other than those explicitly chosen by the microblock author. Also, the ACE approach cannot be used to perform dynamic software reconfiguration as it takes no account of the *integrity* of a running configuration: if a component is replaced, a neighbouring component will inevitably fail as components expect to interact directly. *Teja NP* [13] is another commercial product targeted at the IXP1200, although it also runs on the IBM PowerNP series which is very similar architecturally (at least in terms of our classification scheme) to the IXP1200. Rather than offer an abstract programming model like *MicroACE*, *Teja* focuses on the provision of an integrated tool chain and development environment. Although this eases the development of NP software it provides minimal

architectural abstraction and therefore minimal design portability.

Turning to research-derived programming environments, *NetBind* [14] provides the abstraction of a set of packet-processing components that can be bound into a data path. This is done by adopting the convention of a standard entry and exit instruction sequence for microblocks, and offering the capability to dynamically ‘morph’ jump instructions in these sequences so that execution is transferred to the entry point of the microblock to be executed next. This separates the raw functionality of a microblock from the way it is composed with others, and also gives the *NetBind* programmer the ability to dynamically reconfigure compositions of microblocks.

NetBind goes beyond *MicroACE* in supporting flexible composition of microblocks, but it offers no abstraction over the NP’s memory organisation, interconnects, or over different sorts of processors (e.g. the microengines, StrongARM, and workstation host of an IXP1200-based router). It therefore offers no more design portability across different NPs than *MicroACE*.

NP-Click [15] is another component-based programming model for NPs; it is derived from an earlier PC-based software router model called *Click*. Again, *NP-Click* has been primarily targeted at the IXP1200. It is founded on a much richer model of components than *NetBind*. While communication between *NetBind* microblocks takes place over low-level untyped entry and exit points, *Click* components have typed *ports*; and connections between ports can be designated as either ‘push’ or ‘pull’ which provides declarative control over flow of control and threading. In addition, *NP-Click* abstracts, to a degree, over the different memory technologies offered by the IXP1200 by providing keywords such as ‘global’, ‘regional’ or ‘local’ which cause the associated component to be automatically allocated an appropriate memory type. Furthermore, it provides low level abstractions such as *malloc()* and *free()* to facilitate and manage the allocation of NP resources such as microengine LIFO registers.

NP-Click does a much better job of abstracting NP architecture than *NetBind*, but it still falls short of providing a generic approach to NP programming. While it abstracts particular features of the IXP1200, it has no notion of abstracting arbitrary architectures in a principled way, and thereby encouraging design portability and transferable skills across NP types. That is, there is no necessary commonality between the abstractions provided over different architectures (e.g. NPs other than the IXP1200 may not use LIFOs). In addition, *NP-Click* provides no support for dynamic reconfiguration.

VERA [16], [17] is an extensible software router architecture that comprises three layers: the top layer provides the abstraction of a router, the bottom layer abstracts the hardware, and a middle ‘distributed operating system’ layer mediates between the two. The latter layer organises the available packet processors into a hierarchy of levels. Initial classification occurs on a ‘low level’ processor attached to the MAC-port, and if the packet requires further or more complex processing then a ‘higher level’ processor is used.




Fig. 2. Illustration of capsules and caplets

This provides a high degree of abstraction, but it is heavily dependent on the IXP1200 architecture. For example, it is hard to see how this same abstraction of levels could be maintained on the EZChip NP (see section II).

Apart from the work discussed above, additional research has focused on creating toolsets for specific NPs such as C compilers, simulators, debuggers and benchmarkers; some of this work is described in [18], [19], [20]. Like Teja, however, this work focuses on making tools more usable rather than on providing NP-tailored programming models that promote design portability and transferable programming skills.

Finally, the Network Processor Forum [21], a Industry consortium that aims to facilitate and accelerate the development of NP products, is starting to take an interest in NP programming interface standardisation. To date their focus has been on hardware level interoperability, but they have recently announced the formation of a study group that will define a software API for network-computing applications. However, it is not yet envisaged that this API will address low level programming of individual NP products.

III. TOWARDS A GENERIC PROGRAMMING MODEL FOR NETWORK PROCESSORS

A. Overview of the OpenCOM Component Model

A high-level view of our proposed component model, called *OpenCOM* [22], is given in figure 2. This depicts the central concepts of *components* (the filled circles), *capsules* (the outer dotted box), *caplets* (the inner dotted boxes), *interfaces* (the small circles), *receptacles* (the small cups), and *bindings* (the implied association between the adjacent interfaces and receptacles).

- **Components, Capsules and Caplets** *Components* are encapsulated units of functionality and deployment that interact with their environment (i.e. other components) exclusively through interfaces and receptacles. Components carry negligible inherent overhead and can effectively be used in extremely fine grained compositions. Crucially, OpenCOM is a runtime component model meaning that (unlike, say, NP-Click) components can be dynamically deployed at any time during run-time. The locus of deployment is either a *capsule* or a *caplet*. Both of these concepts represents a scope for component deployment; the latter are sub-scopes of the former (they can be nested to arbitrary depth). In principle (if the deployment environment permits), caplets can be created and

destroyed at run-time. Different caplets can also host components written in different languages (e.g. to accommodate interpreted languages like Java; or to accommodate different machine languages where caplets run on different CPU types).

Each capsule offers a simple run-time API for component lifecycle management (i.e. loading components into the capsule and instantiating and destroying them), and interface/receptacle binding (see below). To accomplish loading, the model supports the notion of *plug-in loaders*. New loaders with different behaviours can be added at runtime, and they can be selected according to their particular properties. Examples are given below. Importantly, the loading of components into a capsule can be requested by any component hosted by the capsule no matter which caplet is hosting it (this is referred to as *third-party deployment*).

- **Interfaces and Receptacles** *Interfaces* are units of service provision offered by components; they are expressed in terms of sets of operation signatures and associated datatypes. For language independence, OMG IDL [23] is used as a specification language. As in Click, components can support multiple interfaces: this is useful in recognising separations of concerns (e.g. between base functionality and management). *Receptacles* are ‘anti-interfaces’ used to make explicit the dependencies of components on other components: whereas an interface represents an element of service provision, a receptacle represents a unit of service *requirement*. Receptacles are key to supporting a third-party style of composition (to complement the third-party deployment referred to above): when third-party-deploying a component into a capsule, one knows by looking at the component’s receptacles precisely which other component types must be present to satisfy its dependencies.

- **Bindings** Finally, *bindings* are associations between a single interface and a single receptacle that reside in a common capsule (but not necessarily a common caplet). Similarly to plug-in loaders, OpenCOM also supports a notion of *plug-in binders*. Again, the idea is to give access to an extensible range of binding mechanisms with varying characteristics. See below for examples. As mentioned, the creation of bindings is inherently third-party in nature; it can be performed by any party within the capsule (i.e. not only by the ‘first-party’ components whose interface or receptacle participates in the binding).

B. Applying OpenCOM in NPs

We now consider how the above concepts can be applied in the diverse range of NP types characterised in section II. First, the scoping-related notions of capsules and caplets are useful in distinguishing different processors and types of processors on the NP in a generic manner (cf. the packet processor dimension). For example, in an IXP1200, we might map a single capsule onto the entire NP, and sub-scope individual microengines, and the StrongARM control processor, as caplets. The capsule runtime in such a mapping would reside on the StrongARM where it could run in a standard operating system environment. An alternative mapping could encapsulate all the

microengines within a single caplet. Furthermore, a plug-in loader associated with this caplet could perform intelligent load balancing of components across microengines, thus providing a higher level of abstraction than the first alternative. The notion of caplets is also useful in isolating untrusted code, which is important in active networking environments. For example, a Java sandbox could be isolated as a caplet.

The IXP1200 is situated towards the ‘centralised’ end of the control dimension defined in section II-A. In an NP with less centralisation, such as the Motorola C-Port or the EZChip, the capsule abstraction could span both the NP itself and its hosting workstation. In this case, the capsule runtime would execute on the host. Alternatively, the capsule abstraction could be restricted to the NP itself, and the capsule runtime could execute on one of the general packet processors, if present. This would be possible in principle on the Motorola C-Port, but not on the EZChip which has no general purpose processors.

The pluggable loader concept is closely associated with capsules/caplets. Typically, at least one loader will be provided for each type of caplet, and each will know how to load components into the hardware (and/ or language) environment underlying its particular caplet type. For example, on the IXP1200 mapping referred to above, there would be (at least) one loader for the StrongARM caplet and another for the microengine caplets. Importantly, the OpenCOM API allows selective transparency in the use of loaders. If full loader-selection transparency is desired, one can issue a call such as *load(component_c1, caplet_1)* which will deduce an appropriate loader type from meta-data attached to *component_c1*, and use this to load the component into the designated caplet. This essentially masks the fact that different components may be implemented in different machine languages. Even more transparency can be requested by issuing a call of the form *load(component_c1)* which causes the runtime to load *component_c1* into a default capsule using a default loader. Alternatively, one can opt for complete control and zero transparency by issuing a call of the form *load(component_c1, caplet_1, loader_3)*.

The pluggable binder concept is equally central to the component model’s abstraction power. In this case, the abstraction is over the interconnect dimension. For example, on the IXP1200 we can encapsulate the NetBind binding mechanism (see section II-B) as a plug-in binder that is competent to bind components that have been loaded into a common caplet that represents a single underlying microengine. But equally well, we can provide a binder that is competent to bind components on different microengines (e.g. based on a shared memory or a next-neighbour register mechanism), or even between components on a microengine and components on the StrongARM. Again, the use of plug-in binders is selectively transparent. If we don’t know or care in which caplets our two target components are located, we can say *bind(interface_1, receptacle_15)* and an appropriate loader will be selected according to location-related meta-data attached to the components that own the specified interface and receptacle. On the

other hand, if it is important to select a particular mechanism, we can say *bind(interface_1, receptacle_15, loader_4)*. And so on.

Note that the abstract model of binding provided by the pluggable binder framework is consistent across all types of NP regardless of the nature and diversity of the interconnects between packet processors. For example, it can uniformly accommodate the fixed hardware channels supported by the pipeline-oriented EZChip, or the bus and shared memory interconnects of the Motorola C-Port, in just the same way as the various mechanisms supported by the IXP1200. Of course, different NP architectures may impose constraints on the form of possible bindings. For example, it would not be straightforward to directly bind components on non-adjacent processors on the EZChip NP; although even here it would be possible (if perhaps undesirable) to provide a plug-in binder that implemented this type of binding by transparently instantiating a forwarder on the intermediary processor(s).

The component concept alone is capable of providing considerable abstraction power in terms of accommodating dedicated non-programmable processors such as those provided by the Motorola C-Port. These processors can be accommodated by representing them with a ‘dummy’ component and an associated special plug-in loader and binder pair. Loading the component and binding it to the client component has the effect of making the service provided by the dedicated processor (e.g. table lookup) look as if it were a normal software component.

A final crucial property of the component model is its radically third-party nature in terms of loading and binding. Thanks to this, a component on an IXP1200 microengine can load and bind two components on the StrongARM control processor, or even on the host workstation, if that comes within the scope of the capsule.

Note that in this paper we omit, for lack of space, any discussion of the important OpenCOM notion of *component frameworks* which are used to support safe dynamic software reconfiguration. Information on this topic is available in the literature [22].

IV. CASE STUDY: OPENCOM ON THE INTEL IXP1200

For the past year we have been working to deploy and evaluate the OpenCOM component model on the Intel IXP1200. The IXP1200 was selected because of its open and well documented architecture, and because it is a richly-featured NP in terms of the dimensions presented in section II-A.

To generate useful components with which to populate the implementation, we have taken as our starting point various modules (e.g. classifiers, forwarders, schedulers etc.) provided by the NetBind project [14] from Columbia University. We have transformed these bare modules, which are written in C or assembler, for both the StrongARM CPU and the microengines, into standard OpenCOM components by attaching appropriate meta-data (e.g. IDL interfaces, and loader and binder attributes) to produce standardly-packaged and deployable units.

The mapping we currently employ of OpenCOM capsules and caplets to the IXP1200 involves a single capsule that encompasses both the NP and the host workstation, and contains separate caplets for: the host workstation (actually, a single Linux process on the workstation); the StrongARM (again, a single Linux process); and each of the six microengines. The OpenCOM runtime runs in the StrongARM caplet; all the other caplets are ‘slaves’ of this ‘central’ runtime and incur only minimal memory overheads (see below). The memory footprint of the central runtime itself is of the order of 300Kb, although we believe that there is considerable scope for reducing this.

The central runtime in the StrongARM caplet communicates with the other caplets by means of so-called *caplet channels*. The role of these is to bootstrap plug-in loaders and binders associated with non-central caplets, and to support communication between their two parts: such loaders/ binders are implemented as a ‘delegator’ part that resides in the central StrongARM caplet, and a (minimal) ‘delegate’ part that resides in the other caplet. As examples, we now briefly describe example loader and binder plug-ins that are associated with the microengine caplets.

The *microengine loader plug-in* is of interest in that it provides the illusion of dynamic loading despite the fact that the microengine hardware only allows modification of its instruction store when the CPU is stopped [12]. The basic capability provided by the microengine hardware is to stop the microengine, read/ write arbitrary instruction store locations, and then restart it at a hard-wired address. To achieve transparent dynamic loading it is therefore necessary for the loader to not only load the new component but also to patch the (hard-wired) restart address so that subsequent execution resumes at the point it left off. The loader also has the ability to autonomously move code around within the instruction store to avoid memory fragmentation as components are loaded and unloaded. The loader is also of interest in that it constrains the form of OpenCOM components it is willing to load. The general notion of particular loaders somehow restricting the components they can work with is a general and powerful pattern in OpenCOM. In the present case, the IDL interfaces of loaded components are restricted to supporting operations that accept and return a single integer. This restriction, which is enforced by inspecting the component’s IDL meta-data at load time, is imposed partly to simplify the design of the associated binder (see below), and partly because the assumed model of component composition on the microengines (borrowed from above-mentioned NetBind work) is that components are bound into a more-or-less linear sequence and cooperatively work on queues of network packets whose addresses are passed as integer arguments.

Our *intra-microengine binder plug-in* is strongly coupled to the loader just described. It builds on the above-mentioned NetBind technique (see section II-B) of creating bindings by ‘morphing’ jump instructions. However, the binder is more complex than the NetBind implementation because, together with the loader, it supports multiple instantiations

of components (unlike NetBind which only supports singleton components). The single argument and return value are passed via a designated register. The necessary entry and exit point information is obtained from IDL meta-data attached to the packaged component, which is transformed from relative offsets to absolute offsets by the loader. It is important to notice, by the way, that the IDL-specified interfaces do *not* incur performance overhead. In fact, the overhead of the binder in calling a null operation with no arguments or return values is only five (1-cycle) instructions. These subsume passing on the stack a pointer to the per-instance state vector of the called component, and the return address. Note that NetBind incurs an overhead of just two 1-cycle jump instructions (for the call and the return). But this is because NetBind does not support multiple instantiations of components. Crucially, however, we could easily retrieve the NetBind performance in the OpenCOM environment simply by implementing and installing a new binder plug-in that understands components that observe the NetBind calling convention and (therefore) does not support multiple component instantiation. The essential point is that OpenCOM’s plug-in architecture enables us to support any appropriate trade-off. More generally, it is important to observe that the performance of the OpenCOM programming model as a whole is almost entirely dependent on the performance of the binding mechanisms used. Almost all the value-added features of OpenCOM are confined to the central runtime and do not ‘get in the way’ when components communicate with each other on the NP’s fast path.

Apart from the microengine loader and binder discussed above, we are currently implementing loaders that load components into StrongARM and host workstation caplets; and binders that bind components across any pairwise combinations of the three caplet types. Bindings between the microengines and the other two caplet types are considerably more complex than intra- and inter-microengine bindings as they require stubs and skeletons to map the parameter and return value to a bus packet. To minimise memory overhead, the microengine-side stubs/ skeletons can be hand coded rather than generated automatically from the IDL specification.

V. CONCLUSIONS AND FUTURE WORK

In this paper we have characterised the design space of NPs and proposed a component-based programming model that, we have argued, can be applied generally within this design space. The component model, mainly through its plug-in loaders and binders and its associated notion of caplets, provides a high degree of design portability and potential for skill transfer. We have also demonstrated how plug-in loaders and binders can exploit NP-specific features to provide both high performance (for example, our microengine binder incurs comparable overheads to NetBind on the IXP1200), and value-added behavior (for example, our microengine loader/ binder supports multiple instantiations of components and transparently optimises instruction store use as components are dynamically loaded/ instantiated/ destroyed).

Most importantly, we have argued that our abstractions are *generally applicable*. NP-Click also abstracts NP-specific features - e.g. it provides an API to manage and allocate microengine LIFO resources on the IXP1200. But this API would make no sense on an NP that did not support LIFOs. The OpenCOM approach would be to provide a plug in binder (a generic abstraction) that internally uses, manages and allocates LIFOs (if present) to build a reusable binder plug-in.

OpenCOM also supports run-time reconfiguration. In this paper we have discussed the basic mechanisms behind this (i.e. receptacles, and the ability to bind and unbind components at runtime). But we have not elaborated on OpenCOM's approach to managing integrity, consistency, safety and security when performing reconfiguration operations. As mentioned, we rely on the notion of *component frameworks* [22] to support this. We have already explored the provision of component frameworks in other domains in which we have applied OpenCOM (e.g. Middleware [24]); one aspect of our future work will be to further explore this interesting and demanding area in the NP domain.

The main thrust of our future work, however, will be to explore the use of OpenCOM in other NP environments. We are already investigating the more advanced IXP2400 from Intel and the IBM PowerNP; but we would also like to provide further evidence for the generality of our approach by looking in more detail at NPs elsewhere in the design space outlined in this paper.

REFERENCES

- [1] L. Ruf, R. Pletka, P. Erni, and P. Droz. Towards High-Performance Active Networking. *Proceedings of the Fifth Annual International Working Conference on Active Networks (IWAN 2003)*, December 2003. Kyoto, Japan.
- [2] C. Kulkarni, M. Gries, C. Sauer, and K. Keutzer. Programming Challenges in Network Processor Deployment. In *Int. Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES)*, October 2003.
- [3] Agere Systems Proprietary. The Challenge for Next Generation Network Processors. April 2001.
- [4] M. Tsai, C. Kulkarni, C. Sauer, N. Shah, and K. Keutzer. A benchmarking methodology for network processors. In *1st Workshop on network processors along with HPCA 2002*, February 2002.
- [5] P. G. Paulin. Network Processors: A Perspective on Market Requirements, Processor Architectures and Embedded S/W Tools. *STMicroelectronics*, 2001.
- [6] M. Adiletta et al. The Next Generation of Intel Network Processors. *Intel Technology Journal, Volume 6, issue 3*, August 15 2002.
- [7] Radisys Corporation. IXP1200 White Paper: Using the Intel IXP1200 Network Processor to Optimize Packet-Processing Application Development, 2001.
- [8] Motorola Research. Architecture guide, C-5e/C-3e Network Processor, Silicon Revision B0, 2003.
- [9] EZchip Technologies. Network Processor Designs for Next-Generation Networking Equipment White Paper, 2003.
- [10] J. Allen et al. IBM PowerNP Network Processor: Hardware Software and Applications. *IBM Journal of Research and Development*, 47(2/3):177-194, March/May 2003.
- [11] Cisco Systems Inc. Parallel Express Forwarding on the Cisco 10000 Series, 2003.
- [12] Intel Press. MicroACE, Design Document, revision 1.0. *Intel Press, Intel Corporation*, 2001.
- [13] A. Deshpande, K. Crozier, and M. Baines. The Teja Software Platform for Network Processors.
- [14] A. T. Campbell, M. E. Kounavis, D. A. Villela, J. B. Vicente, H. G. de Meer, K. Miki, and K. S. Kalaichelvan. NetBind: A Binding Tool for Constructing Data Paths in Network Processor-based Routers. In *5th IEEE International Conference on Open Architectures (OPENARCH'02)*, June 2002.
- [15] N. Shah, W. Plishker, and K. Keutzer. NP-Click: A Programming Model for the Intel IXP1200. In *2nd Workshop on Network Processors (NP-2) at the 9th International Symposium on High Performance Computer Architecture (HPCA-9)*, Anaheim, CA, February 2003.
- [16] S. Karlin and L. Peterson. VERA: An Extensible Router Architecture. In *4th International Conference on Open Architectures and Network Programming (OPENARCH)*, April 2001.
- [17] T. Spalink, S. Karlin, L. Peterson, and Y. Gottlieb. Building a Robust Software-Based Router Using Network Processors, Oct 2001.
- [18] J. Wagner and R. Leupers. C Compiler Design for an Industrial Network Processor. *Proceedings of the 2001 ACM SIGPLAN Workshop on Optimization of Middleware and Distributed Systems*, 2001.
- [19] G. Memik, W. H. Mangione-Smith, and W. Hu. NetBench: A Benchmarking Suite for Network Processors. *ICCAD*, 2001.
- [20] M. Gries, C. Kulkarni, C. Sauer, and K. Keutzer. Comparing Analytical Modeling with Simulation for Network Processors: A Case Study. *Design, Automation, and Test in Europe (DATE)*, Munich, Germany, March, 2003.
- [21] Working Group, Network Processing Forum. Network Processing Forum Background, Oct 2002. <http://www.npforum.org/>.
- [22] G. Coulson, G. Blair, D. Hutchison, A. Joolia, K. Lee, J. Ueyama, A. Gomes, and Y. Ye. NETKIT: A Software Component-Based Approach to Programmable Networking. In *ACM SIGCOMM Computer Communication Review*, volume 33, No 5, October 2003.
- [23] Object Management Group, Inc. CORBA 3.0 - IDL Syntax and Semantics chapter, formal/02-06-07.
- [24] G. Coulson, G. Blair, M. Clarke, and N. Parlavantzas. The Design of a Highly Configurable and Reconfigurable Middleware Platform. *ACM Distributed Computing Journal*, 15(2):109-126, April 2002.